

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, Illinois 60439

Manifold Sampling for L1 Nonconvex Optimization¹

Jeffrey Larson, Matt Menickelly, and Stefan M. Wild

Mathematics and Computer Science Division

Preprint ANL/MCS-P5392-0915

September 2015

¹This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research under contract number DE-AC02-06CH11357.

MANIFOLD SAMPLING FOR L1 NONCONVEX OPTIMIZATION

JEFFREY LARSON*, MATT MENICKELLY†*, AND STEFAN M. WILD*

Abstract. We present a new algorithm, called manifold sampling, for the unconstrained minimization of a nonsmooth composite function $h \circ F$. By classifying points in the domain of the nonsmooth function h into what we call manifolds, we adapt search directions within a trust-region framework based on knowledge of manifolds intersecting the current trust region. We motivate this idea through a study of ℓ_1 functions, where the classification into manifolds using zero-order information about the constituent functions F_i is trivial, and give an explicit statement of a manifold sampling algorithm in that case. We prove that all cluster points of iterates generated by this algorithm are stationary in the Clarke sense. We prove a similar result for a stochastic variant of the algorithm; interestingly, the result is deterministic (not almost sure). Additionally, our algorithm can accept iterates that are points of nondifferentiability and requires only an approximation of gradients of F at the trust-region center. Numerical results presented for several variants of the algorithm show that using manifold information from additional points near the current iterate can improve practical performance. The best variants are also shown to be competitive, particularly in terms of robustness, with other nonsmooth solvers.

Key words. Composite Nonsmooth Optimization, Gradient Sampling, Derivative-Free Optimization

AMS subject classifications. 90C56, 49J52

1. Introduction. This paper addresses the unconstrained optimization problem $\min \{f(x) : x \in \mathbb{R}^n\}$ when f is of the form

$$f(x) = \sum_{i=1}^r |F_i(x)| = \|F(x)\|_1 \quad (1.1)$$

and the function $F : \mathbb{R}^n \rightarrow \mathbb{R}^r$ is sufficiently smooth, as formalized in the following assumption.

ASSUMPTION 1. *The function f is of the form (1.1), each F_i is continuously differentiable, and each ∇F_i is Lipschitz continuous with Lipschitz constant L_i ; define $L = \sum_{i=1}^r L_i$.*

Minimizing the function (1.1) is a special case of more general composite nonsmooth optimization

$$\min \{f(x) = g(x) + h(F(x)) : x \in \mathbb{R}^n\}, \quad (1.2)$$

where g and F are smooth but h is nonsmooth. We focus on the objective function (1.1) in order to succinctly introduce, analyze, and empirically study a general algorithmic framework. Although the form of h studied here is convex, our framework does not require this.

Furthermore, this framework—which we refer to as *manifold sampling*—does not require the availability of the Jacobian ∇F . As a result, manifold sampling is applicable both when inexact values for $\nabla F(x)$ are available and in the derivative-free case, when only the values $F(x)$ are available. In Section 2 we motivate the use of the term manifold in the context of functions of the form (1.1) and show that these manifolds

†Lehigh University, Department of Industrial and Systems Engineering. H.S. Mohler Laboratory, 200 West Packer Avenue, Bethlehem, PA 18015. mjm412@lehigh.edu

*Argonne National Laboratory, Mathematics and Computer Science Division. 9700 South Cass Avenue, Argonne, IL 60439. {jmlarson,wild}@anl.gov

can be determined without using Jacobian information. We also review the literature in composite nonsmooth optimization.

Section 3 introduces a manifold sampling algorithm. The algorithm uses a smooth model M of the mapping F and proceeds like a traditional trust-region algorithm until it encounters an area of possible nondifferentiability. In the case of the function (1.1), a signal for potential nondifferentiability is directly obtained from the signs of the r component functions at the current iterate. We also propose different mechanisms for incorporating such sign information from the current iterate and nearby points.

Under minimal assumptions, in Section 4 we show that our algorithm's smooth, local model of the function f ,

$$m^f(x^k + s) \approx f(x^k) + \left\langle s, \mathbf{proj}\left(0, \nabla M(x^k) \partial h(F(x^k))\right) \right\rangle, \quad (1.3)$$

generates descent directions at the current point x^k . Furthermore, we prove that all cluster points of the algorithm are Clarke stationary points. We show in Section 5 that convergence again holds when using sign information from stochastically generated points.

This stochastic sampling proves to be beneficial in practice, as our numerical tests in Section 6 demonstrate. Our experiments also underscore the relative robustness of four variants of the proposed manifold sampling algorithm. The tested deterministic variants include one that performs efficiently when function evaluations occur sequentially as well as variants that can exploit concurrent function evaluations.

2. Background. We now introduce notation and provide context for our manifold sampling algorithm. We follow the convention throughout this paper that finite sets are denoted by capital Roman letters while (possibly) infinite sets are denoted by capital calligraphic letters.

2.1. Nonsmooth Optimization Preliminaries. We define the set of points at which a function f is differentiable by $\mathcal{D} \subseteq \mathbb{R}^n$ and its complement by \mathcal{D}^c . In smooth optimization, a first-order necessary condition for x to be a local minimum of f is that $\nabla f(x) = 0$. In nonsmooth optimization, if $x \in \mathcal{D}^c$, then one needs a more generalized first-order necessary condition; we achieve this with the generalized gradient ∂f , which is a set-valued function referred to as the *Clarke subdifferential*.

The *Clarke directional derivative* at x in the direction d is given by

$$f^\circ(x; d) = \limsup_{y \rightarrow x, t \downarrow 0} \frac{f(y + td) - f(y)}{t}. \quad (2.1)$$

The Clarke subdifferential at x is the set of linear support functions of $f^\circ(x; d)$ when viewed as a function of $d \in \mathbb{R}^n$:

$$\partial f(x) = \{v \in \mathbb{R}^n : f^\circ(x; d) \geq \langle v, d \rangle \text{ for all } d \in \mathbb{R}^n\}. \quad (2.2)$$

By using the definitions (2.1) and (2.2), one can show (see, e.g., [5, Proposition 2.3.2]) that if f is locally Lipschitz near x and attains a local minimum or maximum at x , then $0 \in \partial f(x)$. Thus, $0 \in \partial f(x)$ can be seen as a nonsmooth analogue of the first-order necessary condition $\nabla f(x) = 0$. This condition, $0 \in \partial f(x)$, is referred to as *Clarke stationarity*.

Furthermore, as a consequence of Rademacher's theorem, if f is locally Lipschitz, then \mathcal{D}^c is a set of Lebesgue measure zero in \mathbb{R}^n . In such cases, an equivalent definition

(see, e.g., [5, Theorem 2.5.1]) of the Clarke subdifferential is

$$\partial f(x) = \mathbf{co} \left\{ \lim_{y^j \rightarrow x} \nabla f(y^j) : \lim_{j \rightarrow \infty} y^j = x, \{y^j : j \geq 1\} \subset \mathcal{D} \right\}, \quad (2.3)$$

where \mathbf{co} denotes the convex hull of a set. Equation (2.3) says that $\partial f(x)$ is the convex hull of all limits of gradients of f at differentiable points in an arbitrarily small neighborhood about x .

2.2. Manifolds of (1.1). The signs of the component functions F_i play a critical role in our focus on functions satisfying Assumption 1. We let sgn be the scalar function that returns the values 1, -1 , and 0 for positive, negative, and null arguments, respectively, and we define $\mathbf{sign} : \mathbb{R}^n \rightarrow \{-1, 0, 1\}^r$ by

$$\mathbf{sign}(x) = \begin{bmatrix} \text{sgn}(F_1(x)) & \text{sgn}(F_2(x)) & \cdots & \text{sgn}(F_r(x)) \end{bmatrix}^T.$$

We say that $\mathbf{sign}(x)$ returns the *sign pattern* of a point x . There exist 3^r possible sign patterns for any $x \in \mathbb{R}^n$; by indexing these possible sign patterns, we define $\mathbf{pat}^q \in \{-1, 0, 1\}^r$ for each $q \in \{1, \dots, 3^r\}$.

For any $x \in \mathbb{R}^n$, its *manifold* $\mathcal{M}(x)$ is the maximal topologically connected set satisfying

$$\mathcal{M}(x) = \{y \in \mathbb{R}^n : \mathbf{sign}(y) = \mathbf{sign}(x)\}.$$

We define the union of all manifolds with the same sign pattern \mathbf{pat}^q as a *manifold set* and denote it by $\mathcal{M}^q = \bigcup_{\{x : \mathbf{sign}(x) = \mathbf{pat}^q\}} \mathcal{M}(x)$. Letting $\mathcal{B}(x, \epsilon) = \{x : \|x\| \leq \epsilon\}$

denote the ball of radius ϵ around a point x , we say that a manifold set \mathcal{M}^q is *active* at x provided that $\mathcal{M}^q \cap \mathcal{B}(x, \epsilon) \neq \emptyset$ for all $\epsilon > 0$.

We note that if Assumption 1 is satisfied, then the function f is locally Lipschitz. By using chain rule results (dependent on regularity conditions; see, e.g., [5, Definition 2.3.4]) that are ensured by Assumption 1, the subdifferential of f at a point x is

$$\partial f(x) = \sum_{i: F_i(x) \neq 0} \text{sgn}(F_i(x)) \nabla F_i(x) + \sum_{i: F_i(x) = 0} \mathbf{co} \{-\nabla F_i(x), \nabla F_i(x)\}, \quad (2.4)$$

where addition is understood to be setwise. Consequently, f is differentiable at x if and only if $F_i(x) = 0$ implies $\nabla F_i(x) = 0$ for $i \in \{1, \dots, r\}$. This observation motivates our definition of the *nondifferentiable set of F_i* , \mathcal{D}_i^c , given by

$$\mathcal{D}_i^c = \{x \in \mathbb{R}^n : F_i(x) = 0 \text{ and } \nabla F_i(x) \neq 0\}. \quad (2.5)$$

Note that for a function satisfying Assumption 1, $\mathcal{D}^c = \bigcup_{i=1}^r \mathcal{D}_i^c$. Our algorithmic framework in Section 3 is predicated on a relaxation of \mathcal{D}_i^c that does not require the (exact) derivative $\nabla F_i(x)$.

2.3. Related Work. The analogue to steepest descent for nonsmooth optimization involves steps along the negative of the minimum norm element of the subdifferential, $-\mathbf{proj}(0, \partial f(x^k))$. The algorithm that we propose is related to bundle methods [15] in that the subdifferential in this step is approximated by a finite set of generators.

A class of methods for nonsmooth optimization related to both bundle methods and our proposed algorithm is that of gradient sampling. Such methods exploit situations when the underlying nonsmooth function is differentiable almost everywhere by using local gradient information around a current iterate to build a stabilized descent direction [4]. For example, under weak regularity conditions, the closure of the set $\text{co}\{\nabla f(y) : y \in \mathcal{B}(x^k, \Delta) \cap \mathcal{D}\}$ is a superset of $\partial f(x^k)$. Gradient sampling methods employ a finite sample set $\{y^1, \dots, y^p\} \subset \mathcal{B}(x^k, \Delta) \cap \mathcal{D}$ and approximate $\partial f(x^k)$ by $\text{co}\{\nabla f(y^1), \dots, \nabla f(y^p)\}$. The negative of the minimum norm element of this latter set is used as the search direction.

Constructing this set is more difficult when ∇f is unavailable. A nonderivative version of the gradient sampling algorithm is shown in [17] to converge almost surely to a first-order stationary point. However, the analysis depends on the use of a Gupal estimate of Steklov averaged gradients as a gradient approximation. Such an approach requires $2n$ function evaluations to compute each approximate gradient. In effect, a single gradient approximation is as expensive to compute as a central difference gradient approximation, and the approximations must be computed for each point near the current iterate. With this motivation, Hare and Nutini [13] propose an approximate gradient-sampling algorithm that uses standard (e.g., central difference, simplex) gradient approximations to solve finite minimax problems of the composite form $\max_i F_i(x)$. Similar to our focus on the form (1.1), Hare and Nutini [13] model the smooth component functions F_i and use the particular, known structure of their subdifferential $\text{co}\{\nabla F_j(x) : j \in \arg \max_i F_i(x)\}$ within their convergence analysis. Both [13] and [17] employ a line search strategy for globalization. Our method employs a trust-region framework that links the trust-region radius with the norm of a model gradient, which can serve as a stationarity measure.

Other methods also exploit the general structure in composite nonsmooth optimization problems (1.2). When the function h is convex, typical trust-region-based approaches (e.g., [3, 8, 9, 22]) solve the *nonsmooth* subproblem

$$\min \{h(F(x^k) + \langle s, \nabla M(x^k) \rangle) : s \in \mathcal{B}(x^k, \Delta)\}. \quad (2.6)$$

The model M is a Taylor expansion of F when derivatives are available. In this case, Griewank et al. [12] propose building piecewise linear models using ∇F from recent function evaluations. These nonsmooth models are then minimized to find future iterates and ∇F information for subsequent models. When derivatives are unavailable, recent work has used sufficiently accurate models of the Jacobian [10, 11]. We follow a similar approach in our approximation of ∇F but employ a fundamentally different subproblem, locally minimizing smooth models related to (1.3). In contrast to (2.6), our subproblem does not rely on the convexity of h .

There also exist derivative-free methods for nonsmooth optimization unrelated to gradient sampling. For example, a generalization of mesh adaptive direct search [1] can find descent directions for nonsmooth problems by generating an asymptotically dense set of search directions. Similar density requirements exist for general direct search methods [20]. One of the derivative-free methods in [10] employs a smoothing function $f_\mu(x)$ that is parameterized by a smoothing parameter $\mu > 0$ satisfying

$$\lim_{z \rightarrow x, \mu \downarrow 0} f_\mu(z) = f(x),$$

for any $x \in \mathbb{R}^n$. By iteratively driving $\mu \rightarrow 0$ within a trust-region framework, the authors prove convergence to Clarke stationary points and provide convergence rates.

We note that the Steklov averaged gradients in [17] are also essentially smoothing convolution kernels. Our proposed method requires neither a dense set of search directions nor smoothing parameters.

3. Algorithmic Framework. We now outline our algorithm, which samples manifolds (as opposed to gradients or gradient approximations) in order to approximate the subdifferential $\partial f(x^k)$. When x^k is changed, the r component function values $F_i(x^k)$ are computed, immediately yielding $\mathbf{sign}(x^k)$. Then, r models, m^{F_i} , approximating F_i near x^k are built. Using the value of $\mathbf{sign}(x^k)$, we infer a set of generators, G^k , using the manifolds that are potentially active at x^k . The set $\mathbf{co}(G^k)$, detailed in Section 3.2, is used as an approximation to $\partial f(x^k)$.

We let g^k denote the minimum norm element of $\mathbf{co}(G^k)$,

$$g^k = \mathbf{proj}(0, \mathbf{co}(G^k)), \quad (3.1)$$

which can be achieved by solving the quadratic optimization problem

$$\min \left\{ \frac{1}{2} \lambda^T G^T G \lambda : e^T \lambda = 1, \lambda \geq 0 \right\}, \quad (3.2)$$

where the columns of G are the generators in G^k . A solution λ^* to (3.2) is a set of weights on the subgradient approximations that minimize $\|G\lambda\|^2$. That is, $g^k = G\lambda^*$.

The minimum norm element g^k will define the first-order term of a smooth *master model*, $m_k^f : \mathbb{R}^n \rightarrow \mathbb{R}$ and thereby the trust-region subproblem

$$\min \left\{ m_k^f(x) : x \in \mathcal{B}(x^k, \Delta_k) \right\}, \quad (3.3)$$

which has the closed-form solution $s^k = -\frac{\Delta_k}{\|g^k\|} g^k$ when m_k^f is linear.

Provided the solution $x^k + s^k$ of (3.3) belongs to a manifold that was involved in the construction of g^k , we apply a ratio test to determine the successfulness of the proposed step, as in a standard trust-region method. If this manifold is not contributing a generator in G^k , we augment G^k and construct a new g^k . Since there are finitely many manifolds (at most 3^r in the case of (1.1)), this process will terminate.

Our ratio test quantity

$$\rho_k = \frac{\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k + s^k)}{\langle -s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle} \quad (3.4)$$

differs from the usual ratio of actual reduction to predicted reduction only in that it considers both the function and model from the perspective of the manifold of $x^k + s^k$. In particular, our numerator is more conservative than the actual reduction since

$$\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k) = \langle F(x^k), \mathbf{sign}(x^k + s^k) - \mathbf{sign}(x^k) \rangle \leq 0,$$

where the inequality holds because $F_i(x^k)[\mathbf{sign}(x^k)]_i = |F_i(x^k)| \geq t_i F_i(x^k)$ for any i and for all $t_i \in \{-1, 0, 1\}$. When m_k^f is linear, the relationship between the denominator of our ρ_k and the model reduction is

$$\begin{aligned} & \langle -s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle - (m_k^f(x^k) - m^f(x^k + s^k)) \\ &= \langle -s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) - g^k \rangle = \frac{\Delta_k}{\|g^k\|} \langle g^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) - g^k \rangle, \end{aligned}$$

Algorithm 1: Manifold sampling.

```

1 Set parameters  $\eta_1, \eta_2, \gamma_{\text{dec}}$ , and  $\gamma_{\text{inc}}$ 
2 Choose initial iterate  $x^0$  and trust-region radius  $\Delta_0 > 0$ ; set  $k = 0$ 
3 while true do
4   For each  $F_i$ , build a model  $m^{F_i}$  that is fully linear in  $\mathcal{B}(x^k, \Delta_k)$ 
5   Build a set of generators  $G^k$ 
6   while true do
7     Compute  $g^k$  using (3.1), and build master model  $m_k^f$ 
8     Solve (3.3) to obtain  $x^k + s^k$ , and evaluate  $f(x^k + s^k)$ 
9     if  $\nabla M(x^k) \text{sign}(x^k + s^k) \in G^k$  then
10      | break; go to Line 13
11    else
12      |  $G^k \leftarrow G^k \cup \nabla M(x^k) \text{sign}(x^k + s^k)$ 
13  if  $\Delta_k < \eta_2 \|g^k\|$  (acceptable iteration) then
14    | Update  $\rho_k$  through (3.4)
15    | if  $\rho_k > \eta_1$  (successful iteration) then
16      |  $x^{k+1} \leftarrow x^k + s^k, \Delta_{k+1} \leftarrow \gamma_{\text{inc}} \Delta_k$ 
17    | else
18      |  $x^{k+1} \leftarrow x^k, \Delta_{k+1} \leftarrow \gamma_{\text{dec}} \Delta_k$ 
19  else
20    |  $x^{k+1} \leftarrow x^k, \Delta_{k+1} \leftarrow \gamma_{\text{dec}} \Delta_k$ 
21   $k \leftarrow k + 1$ 

```

which is nonpositive on acceptable iterations (as shown in Lemma 3).

We now state our framework in Algorithm 1 and assumptions on the algorithmic parameters in Assumption 2.

ASSUMPTION 2. $\eta_1 \in (0, 1)$, $\eta_2 \in (0, \infty)$, $\gamma_{\text{dec}} \in (0, 1)$, and $\gamma_{\text{inc}} \geq 1$.

3.1. Model Quality. In order to prove convergence of our algorithm, the models m^{F_i} must sufficiently approximate F_i in a neighborhood of x^k . We require that the models be *fully linear* in the trust region, a notion formalized in the following assumption.

ASSUMPTION 3. For $i \in \{1, \dots, r\}$, let m^{F_i} denote a model with Lipschitz continuous gradient intended to approximate F_i on some $\mathcal{B}(x, \Delta)$. For each $i \in \{1, \dots, r\}$, for all $x \in \mathbb{R}^n$, and for all $\Delta > 0$, there exist constants $\kappa_{i,\text{ef}}$ and $\kappa_{i,\text{eg}}$, independent of x and Δ , so that

$$\begin{aligned} |F_i(x + s) - m^{F_i}(x + s)| &\leq \kappa_{i,\text{ef}} \Delta^2 \quad \forall s \in \mathcal{B}(0, \Delta) \\ \|\nabla F_i(x + s) - \nabla m^{F_i}(x + s)\| &\leq \kappa_{i,\text{eg}} \Delta \quad \forall s \in \mathcal{B}(0, \Delta). \end{aligned}$$

For these constants, define $\kappa_{\text{f}} = \sum_{i=1}^r \kappa_{i,\text{ef}}$ and $\kappa_{\text{g}} = \sum_{i=1}^r \kappa_{i,\text{eg}}$.

Assumption 3 is nonrestrictive, and one can derive classes of models satisfying the assumption both when ∇F_i is available inexactly and when ∇F_i is unavailable. For instance, in the latter case, the assumption holds when m^{F_i} is a linear model interpolating F_i at a set of sufficiently affinely independent points (see, e.g., [6, Chapter 10]); in this case, $\kappa_{i,\text{ef}}$ and $\kappa_{i,\text{eg}}$ scale with n and the respective Lipschitz constants of m^{F_i} and ∇m^{F_i} .

3.2. Generator Sets. We complete our description of the manifold sampling algorithm by showing how the set of generators G^k is built so that $\text{co}(G^k)$ approximates $\partial f(x^k)$. Given the definition of the Clarke subdifferential in (2.3) and the known form of the subdifferential in (2.4), we know that the extreme points of $\partial f(x)$ must be the limits of sequences of gradients at differentiable points from manifolds that are active at x . Therefore, the extreme points of $\partial f(x)$ are a subset of $\nabla F(x)\text{pat}^q$ over $q \in \{1, \dots, 3^r\}$.

For any $q \in \{1, \dots, 3^r\}$, an approximation

$$\nabla M(x) = [\nabla m^{F_1}(x), \dots, \nabla m^{F_r}(x)]$$

of the Jacobian ∇F induces an approximation $\nabla M(x)\text{pat}^q$ to $\nabla F(x)\text{pat}^q$. Since these models may be inexact, we relax the dependence on ∇F_i in (2.5) and thereby include terms of the form $\text{co}\{-\nabla m^{F_i}(x), \nabla m^{F_i}(x)\}$ whenever $F_i(x) = 0$.

This is the motivation for our first procedure for forming the generator set G^k . Algorithm 2 initializes G^k with $\nabla M(x^k)\text{sign}(x^k)$ and then triples the size of G^k for each i satisfying $\text{sgn}(F_i(x^k)) = 0$. Our analysis of Algorithm 1 will show that this strategy allows us to obtain an approximation of the subdifferential $\partial f(x^k)$.

Algorithm 2: Forming generator set G^k using possibly active manifolds at x^k .

```

1 Input:  $x^k$  and  $\nabla M(x^k)$ 
2  $G^k \leftarrow \{\nabla M(x^k)\text{sign}(x^k)\}$ 
3 for  $i = 1, \dots, r$  do
4   if  $\text{sgn}(F_i(x^k)) = 0$  then
5      $G^k \leftarrow G^k \cup \{G^k + \nabla m^{F_i}(x^k)\} \cup \{G^k - \nabla m^{F_i}(x^k)\}$ 

```

We propose a second approach for constructing G^k that uses sign patterns of points near x^k and not just $\text{sign}(x^k)$. Although this approach is inspired by gradient sampling, we note that we are not approximating the gradient at any point other than x^k . We naturally extend our definition of active manifolds by saying that a manifold set \mathcal{M}^q is active in a set \mathcal{S} provided there exists $x \in \mathcal{S}$ such that \mathcal{M}^q is active at x . We denote such a *sample set* at iteration k by $Y(x^k, \Delta_k) \subset \mathcal{B}(x^k, \Delta_k)$. This set can come, for example, from the set of all points previously evaluated by the algorithm that lie within a distance Δ_k of x^k .

We can now state Algorithm 3, which constructs generators based on the set of manifolds active in $Y(x^k, \Delta_k)$. Intuitively, this additional manifold information obtained from sampling can “warn” the current iterate about sudden changes in gradient behavior that may occur within the current trust region.

Algorithm 3: Forming generator set G^k using possibly active manifolds in $Y(x^k, \Delta_k)$.

```

1 Input:  $x^k$ ,  $\nabla M(x^k)$ , and  $Y(x^k, \Delta_k) = \{x^k, y^2, \dots, y^p\}$ 
2 Initialize  $G^k$  using Algorithm 2( $x^k, \nabla M(x^k)$ )
3 for  $j = 2, \dots, p$  do
4    $G^k \leftarrow G^k \cup \nabla M(x^k)\text{sign}(y^j)$ 

```

Other reasonable approaches for constructing G^k exist. For our analysis, a requirement for G^k is given in Assumption 4.

ASSUMPTION 4. *Given x^k and $\Delta_k > 0$, the constructed set G^k satisfies*

$$G^k \subseteq \left\{ \nabla M(x^k) \mathbf{sign}(x^k) + \sum_{i: [\mathbf{sign}(x^k)]_i = 0} t_i \nabla m^{F_i}(x^k) : t \in \{-1, 0, 1\}^r \right\} \subseteq G^k, \\ G^k \subseteq \left\{ \nabla M(x^k) \mathbf{sign}(y) + \sum_{i: [\mathbf{sign}(y)]_i = 0} t_i \nabla m^{F_i}(x^k) : t \in \{-1, 0, 1\}^r, y \in \mathcal{B}(x^k; \Delta_k) \right\}.$$

Clearly, a set G^k produced by Algorithm 2 or Algorithm 3 will satisfy Assumption 4. Furthermore, any generator set satisfying Assumption 4 has $|G^k| \leq 3^r$.

4. Analysis. We now analyze the iterates of Algorithm 1.

4.1. Preliminaries. We first show a result linking elements in a set similar to the form of G^k to the subdifferentials of f at nearby points. Subsequent results will establish cases when our construction of the generator set G^k satisfies the suppositions made in the statement of the lemma.

LEMMA 1. *Let Assumptions 1 and 3 hold, and let $x, y \in \mathbb{R}^n$ satisfy $\|x - y\| \leq \Delta_k$. Suppose that $\{\mathbf{pat}^{q_s} : s = 1, \dots, j\} \subseteq \{\mathbf{pat}^{q'_{s'}} : s' = 1, \dots, j'\}$ for*

$$G = \{\nabla M(x) \mathbf{pat}^{q_s} : s = 1, \dots, j\} \text{ and } \partial f(y) = \mathbf{co} \left\{ \nabla F(y) \mathbf{pat}^{q'_{s'}} : s' = 1, \dots, j' \right\}.$$

Then for each $g \in \mathbf{co}(G)$, there exists $v(g) \in \partial f(y)$ satisfying

$$\|g - v(g)\| \leq (\kappa_g + L) \Delta_k, \quad (4.1)$$

where κ_g and L are defined in Assumption 3 and Assumption 1, respectively.

Proof. Let $g \in \mathbf{co}(G)$ be arbitrary. Since $\mathbf{co}(G)$ is finitely generated (and thus compact and convex), g can be expressed as a convex combination of $N \leq n + 1$ of its generators due to Caratheodory's theorem. Without loss of generality (by reordering as necessary), let these generators be the first N . That is, there exist $\lambda_{q_1}, \dots, \lambda_{q_N} \in (0, 1]$ with $\sum_{s=1}^N \lambda_{q_s} = 1$ so that

$$g = \sum_{s=1}^N \lambda_{q_s} \nabla M(x) \mathbf{pat}^{q_s}. \quad (4.2)$$

By supposition, $\nabla F(y) \mathbf{pat}^{q_s} \in \partial f(y)$ for $s = 1, \dots, N$. Since $\partial f(y)$ is convex, we have that $v(g) \in \partial f(y)$, where $v(g)$ is defined as

$$v(g) = \sum_{s=1}^N \lambda_{q_s} \nabla F(y) \mathbf{pat}^{q_s},$$

using the same λ_{q_s} as in (4.2) for $s = 1, \dots, N$. Observe that for each s ,

$$\begin{aligned} \|\nabla M(x) \mathbf{pat}^{q_s} - \nabla F(y) \mathbf{pat}^{q_s}\| &= \|(\nabla M(x) - \nabla F(x) + \nabla F(x) - \nabla F(y)) \mathbf{pat}^{q_s}\| \\ &\leq \|\nabla M(x) - \nabla F(x)\| + \|\nabla F(x) - \nabla F(y)\| \\ &\leq (\kappa_g + L) \Delta_k. \end{aligned}$$

Applying the definitions of g and $v(g)$ yields the expression (4.1). \square

The approximation property in Lemma 1 can be used to motivate the use of the master model gradient in (3.1). Before demonstrating how the smooth master model identifies descent directions, we state a canonical result [2, Proposition 1.1.9], which we refer to throughout as the “Projection Theorem.”

THEOREM 2. *Let G be a nonempty closed convex subset of \mathbb{R}^n , and let z be a vector in \mathbb{R}^n . There exists a unique vector $g^* \in G$ that minimizes $\|z - g\|$, referred to as the projection of z onto G . Moreover, g^* is the projection of z onto G if and only if $\langle g - g^*, g^* - z \rangle \geq 0$ for all $g \in G$.*

4.2. Analysis of Algorithm 1. We now demonstrate that building a master model gradient g^k from the gradients of the models of manifolds that are active in the trust region ensures a successful iteration if Δ_k is sufficiently small.

LEMMA 3. *Let Assumptions 1–3 hold, let m_k^f be linear for all k , and let every component model be of the form $m^{F_i}(x^k + s) = F_i(x^k) + \langle \nabla m^{F_i}(x^k), s \rangle$. If $\Delta_k < \min \left\{ \frac{1-\eta_1}{\kappa_f}, \eta_2 \right\} \|g^k\|$, then iteration k of Algorithm 1 is successful.*

Proof. Notice that the bound on Δ_k is positive by Assumption 2 and that $\Delta_k < \eta_2 \|g^k\|$ ensures that the iteration is acceptable. Since m_k^f is linear, $s^k = -\frac{\Delta_k}{\|g^k\|} g^k$. Thus, the definition of ρ_k in Line 14 of Algorithm 1 yields

$$\begin{aligned} |\rho_k - 1| &= \left| \frac{\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k + s^k)}{\langle -s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle} - 1 \right| \\ &= \left| \frac{\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k + s^k) + \langle s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle}{\frac{\Delta_k}{\|g^k\|} \langle g^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle} \right|. \end{aligned} \quad (4.3)$$

By using the linear form for each component model, the numerator of (4.3) satisfies

$$\begin{aligned} & \left| \langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k + s^k) + \langle s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle \right| \\ &= \left| f(x^k + s^k) - \sum_{i=1}^r m^{F_i}(x^k + s^k) [\mathbf{sign}(x^k + s^k)]_i \right| \\ &= \left| \sum_{i=1}^r (F_i(x^k + s^k) - m^{F_i}(x^k + s^k)) [\mathbf{sign}(x^k + s^k)]_i \right| \\ &\leq \sum_{i=1}^r |F_i(x^k + s^k) - m^{F_i}(x^k + s^k)| \\ &\leq \kappa_f \Delta_k^2, \end{aligned} \quad (4.4)$$

where the last inequality is due to Assumption 3. We can also bound the denominator of (4.3) through

$$\begin{aligned}
& \left| \frac{\Delta_k}{\|g^k\|} \langle g^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle \right| \\
&= \left\langle \nabla M(x^k) \mathbf{sign} \left(x^k - \Delta_k \frac{g^k}{\|g^k\|} \right), \Delta_k \frac{g^k}{\|g^k\|} \right\rangle \\
&= \left\langle \nabla M(x^k) \mathbf{sign} \left(x^k - \Delta_k \frac{g^k}{\|g^k\|} \right) - g^k + g^k, \Delta_k \frac{g^k}{\|g^k\|} \right\rangle \\
&\geq \Delta_k \|g^k\|,
\end{aligned} \tag{4.5}$$

where the last inequality follows from the Projection Theorem and the fact that Algorithm 1 requires that $\nabla M(x^k) \mathbf{sign}(x^k + s^k) \in G^k$ before s^k can be accepted. Using (4.4) and (4.5) in (4.3) yields

$$|\rho_k - 1| \leq \frac{\kappa_f \Delta_k^2}{\Delta_k \|g^k\|} = \frac{\kappa_f \Delta_k}{\|g^k\|} < 1 - \eta_1,$$

where the last inequality is by the supposition on Δ_k . Thus, $\rho_k > \eta_1$, and iteration k is successful. \square

The next result shows that the trust-region radius converges to zero.

LEMMA 4. *Let Assumptions 1 and 2 hold, and let m_k^f be linear for all k . If $\{x^k, \Delta_k\}$ is generated by Algorithm 1, then $\Delta_k \rightarrow 0$.*

Proof. On successful iterations k , the linearity of m_k^f implies that

$$\begin{aligned}
\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle - f(x^k + s^k) &> \eta_1 \langle -s^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle \\
&= \eta_1 \frac{\Delta_k}{\|g^k\|} \langle g^k, \nabla M(x^k) \mathbf{sign}(x^k + s^k) \rangle \\
&\geq \eta_1 \Delta_k \|g^k\| \\
&> \frac{\eta_1}{\eta_2} \Delta_k^2,
\end{aligned}$$

where the last two inequalities follow by (4.5) and the acceptability of all successful steps, respectively. If there are infinitely many successful iterations, let $\{k_j\}$ index them. Notice that on any iteration,

$$\langle F(x^k), \mathbf{sign}(x^k + s^k) \rangle \leq \langle F(x^k), \mathbf{sign}(x^k) \rangle = f(x^k),$$

since, for any i , $F_i(x^k)[\mathbf{sign}(x^k)]_i = |F_i(x^k)| \geq t_i F_i(x^k)$ for all $t_i \in \{-1, 0, 1\}$.

Since f is bounded below by zero by Assumption 1 and $f(x^k)$ is nonincreasing in k , having infinitely many successful iterations implies that

$$\begin{aligned}
\infty &> \sum_{j=0}^{\infty} f(x^{k_j}) - f(x^{k_j} + s^{k_j}) \geq \sum_{j=0}^{\infty} \langle F(x^{k_j}), \mathbf{sign}(x^{k_j} + s^{k_j}) \rangle - f(x^{k_j} + s^{k_j}) \\
&> \sum_{j=0}^{\infty} \frac{\eta_1}{\eta_2} \Delta_{k_j}^2.
\end{aligned}$$

Thus, $\Delta_{k_j} \rightarrow 0$ provided $\{k_j\}$ is an infinite subsequence. Since Δ_k increases by γ_{inc} on successful iterates, for any successful iterate k_i , $\gamma_{\text{inc}} \Delta_{k_i} \geq \Delta_j \geq \Delta_{k_{i+1}}$ for all $k_i < j \leq k_{i+1}$. Therefore, $\Delta_k \rightarrow 0$ if the number of successful iterates is infinite.

If there are only finitely many successes, then there is a last successful iteration $k_{j'}$, and the update rules of the algorithm will monotonically decrease Δ_k on all iterations $k > k_{j'}$.

We have showed that regardless of whether Algorithm 1 has an infinite or a finite number of successful iterations, $\Delta_k \rightarrow 0$. \square

We now show that the norms of the master model gradients are not bounded away from zero.

LEMMA 5. *Let Assumptions 1–3 hold, let m_k^f be linear for all k , and let every component model be of the form $m^{F_i}(x^k + s) = F_i(x^k) + \langle \nabla m^{F_i}(x^k), s \rangle$. If $\{x^k, \Delta_k\}$ is generated by Algorithm 1, then for all $\epsilon > 0$, there exists a $k(\epsilon)$ such that $\|g^{k(\epsilon)}\| \leq \epsilon$. That is, $\liminf_{k \rightarrow \infty} \|g^k\| = 0$.*

Proof. To arrive at a contradiction, suppose that there exist j and $\epsilon > 0$ so that for all $k \geq j$, $\|g^k\| \geq \epsilon$. By Lemma 3, since Algorithm 1 requires that $\nabla M(x^k) \text{sign}(x^k + s^k) \in G^k$ before s^k can possibly be accepted, any iteration satisfying $\Delta_k \leq \min \left\{ \frac{1 - \eta_1}{\kappa_f}, \eta_2 \right\} \|g^k\|$ will be successful. Hence, by the contradiction hypothesis, the first $k \geq j$ satisfying

$$\Delta_k \leq \min \left\{ \frac{1 - \eta_1}{\kappa_f}, \eta_2 \right\} \epsilon \quad (4.6)$$

is guaranteed to be successful and $\Delta_{k+1} = \gamma_{\text{inc}} \Delta_k \geq \Delta_k$. Thus Δ_k does not decrease when $k \geq j$ and Δ_k satisfies (4.6) contradicting Lemma 4. \square

Before showing that every cluster point of $\{x^k\}$ is a Clarke stationary point, we recall basic terms and a proposition from [21].

Motivated by the subdifferential operator $\partial f(x)$, we first formalize the notion of a limit superior of a set mapping (i.e., a mapping that can map a vector to a set of vectors). For a set mapping $D : \mathbb{R}^n \rightarrow \mathbb{R}^n$, the *limit superior* of D as $x \rightarrow \bar{x}$ is defined by the set mapping

$$\limsup_{x \rightarrow \bar{x}} D(x) = \{y : \exists \{x^k : k \geq 1\} \rightarrow \bar{x} \text{ and } \{y^k : k \geq 1\} \rightarrow y \text{ with } y^k \in D(x^k)\}.$$

A set mapping D is said to be *outer semicontinuous at \bar{x}* provided

$$\limsup_{x \rightarrow \bar{x}} D(x) = D(\bar{x}).$$

The following result is given as Proposition 8.7 in [21].

THEOREM 6. *If a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is continuous, then the set mapping $\partial f(x)$ is everywhere outer semicontinuous.*

We now prove the promised result.

THEOREM 7. *Let Assumptions 1–3 hold, let m_k^f be linear for all k , let every component model be of the form $m^{F_i}(x^k + s) = F_i(x^k) + \langle \nabla m^{F_i}(x^k), s \rangle$, and let Assumption 4 hold. If x^* is a cluster point of a sequence $\{x^k\}$ generated by Algorithm 1, then $0 \in \partial f(x^*)$.*

Proof. Suppose first that there are only finitely many successful iterations, with k' being the last.

To establish a contradiction, suppose $0 \notin \partial f(x^{k'})$. By the continuity of each component F_i granted by Assumption 1, there exists $\bar{\Delta}$ so that for all $\Delta \in [0, \bar{\Delta}]$, the manifold sets active in $\mathcal{B}(x^{k'}, \bar{\Delta})$ are precisely the manifold sets active at $x^{k'}$; that is,

$$\left\{ \lim_{y^j \rightarrow x^{k'}} \text{sign}(y^j) : \lim_{j \rightarrow \infty} y^j = x^{k'}, \{y^j : j \geq 1\} \subset \mathcal{M}^q \right\} = \left\{ \text{sign}(y) : y \in \mathcal{B}(x^{k'}, \Delta) \right\}.$$

Since every iteration after k' is assumed unsuccessful, hence decreasing $\Delta_{k'}$ by an additional factor of γ_{dec} in each subsequent iteration, there exists a least $k'' \geq k'$ so that $\Delta_{k''} < \bar{\Delta}$. Therefore, by Assumption 4, $\nabla M(x^k) \mathbf{sign}(x^k + s^k) \in G^k$ holds the first time Line 9 of Algorithm 1 is reached in iteration $k \geq k''$. Consequently, the conditions for Lemma 1 hold; and thus, for each $k \geq k''$, there exists $v(g^k) \in \partial f(x^{k'})$ so that $\|v(g^k) - g^k\| \leq \kappa_g \Delta_k$. By supposition, since $0 \notin \partial f(x^{k'})$, there is a nonzero minimum norm element $v^* \in \partial f(x^{k'})$. We thus conclude the following:

$$\|g^k\| \geq \|v(g^k)\| - (\kappa_g + L)\Delta_k \geq \|v^*\| - (\kappa_g + L)\Delta_k \quad \text{for all } k \geq k''. \quad (4.7)$$

Since every iteration after k'' is unsuccessful, Δ_k will decrease by a factor of γ_{dec} and $x^{k+1} = x^{k'}$ for each $k \geq k''$. Defining the constant $c_1 = \frac{1 - \eta_1}{\kappa_f + (\kappa_g + L)(1 - \eta_1)}$, it must be the case that

$$\Delta_{k_1} < c_1 \|v^*\| \leq \frac{(1 - \eta_1)}{\kappa_f} \|g^{k_1}\|,$$

for some $k_1 \geq k''$ sufficiently large. Similarly using (4.7), we get that

$$\Delta_{k_2} < \frac{\eta_2 \|v^*\|}{1 + \eta_2(\kappa_g + L)} \leq \eta_2 \|g^{k_2}\|,$$

for some $k_2 \geq k''$ sufficiently large. Defining $c_2 = \|v^*\| \min \left\{ c_1, \frac{\eta_2}{1 + \eta_2(\kappa_g + L)} \right\}$, we

have that success is guaranteed within $\left\lceil \log_{\gamma_{\text{dec}}} \frac{c_2}{\Delta_{k''}} \right\rceil$ many iterations after iteration k'' , since decreasing $\Delta_{k''}$ by that many factors of γ_{dec} results in all the sufficient conditions of Lemma 3 being met. This is a contradiction, thus proving the result when there are finitely many successful iterations.

Now suppose there are infinitely many successful iterations, indexed by $\{k_j\}$. Let x^* be a cluster point of $\{x^k : k \geq 0\}$. By the same argument when the number of successful iterations was assumed to be finite, there exists some $\bar{\Delta} > 0$ so that the manifold sets active in $\mathcal{B}(x^*, \bar{\Delta})$ are precisely the manifold sets active at x^* .

Since $\Delta_k \rightarrow 0$ by Lemma 4 and since x^* is a cluster point, there exists an infinite subsequence of $\{k_j\}$, denoted by $\{k_{j'}\}$, so that $\mathcal{B}(x^{k_{j'}}, \Delta_{k_{j'}}) \subset \mathcal{B}(x^*, \bar{\Delta})$ for all j' . On each of these iterations $k_{j'}$, denoting

$$G = \{\nabla M(x^{k_{j'}}) \mathbf{pat}^{q_s} : s = 1, \dots, l\} \quad \text{and} \\ \partial f(x^*) = \mathbf{co} \left\{ \nabla F(x^*) \mathbf{pat}^{q'_{s'}} : s' = 1, \dots, l' \right\},$$

we have from Assumption 4 that $\{\mathbf{pat}^{q_s} : s = 1, \dots, l\} \subseteq \{\mathbf{pat}^{q'_{s'}} : s' = 1, \dots, l'\}$ since $\mathcal{B}(x^{k_{j'}}, \Delta_{k_{j'}}) \subset \mathcal{B}(x^*, \bar{\Delta})$.

Therefore, by Lemma 1 and the reverse triangle inequality, there exists some nonzero $v(g^{k_{j'}}) \in \partial f(x^{k_{j'}})$ satisfying

$$\|v(g^{k_{j'}})\| \leq \|g^{k_{j'}}\| + (\kappa_g + L)\Delta_{k_{j'}},$$

where κ_g and L are defined in Assumption 3 and Assumption 1, respectively.

Since $\Delta_{k_{j'}} \rightarrow 0$ by Lemma 4 and $\|g^{k_{j'}}\| \rightarrow 0$ by Lemma 5, we get by the squeeze theorem that there exists a sequence $v(g^{k_{j'}})$ with $v(g^{k_{j'}}) \in \partial f(x^{k_{j'}})$ for all $\{k_{j'}\}$ and

$\|v(g^{k_{j'}})\| \rightarrow 0$. Since f is continuous, we know that the mapping $\partial f(x)$ is everywhere outer semicontinuous by Theorem 6. Because we have a subsequence of subgradients that converge to zero at any cluster point x^* and since $\partial f(x^*)$ is outer semicontinuous, we have that $0 \in \partial f(x^*)$. \square

4.3. Finite Time Behavior. One would hope that (on acceptable iterations), a small master model gradient norm would signal that necessary conditions for proximity to stationarity are satisfied, analogous to how a small gradient norm serves as such a signal in smooth optimization. Although this is not always so, we provide the exact theoretical conditions under which a statement like this holds for Algorithm 1. This approach emulates a stopping criterion used in [13], which likewise cannot generally be shown to be a necessary condition of proximity to stationarity.

LEMMA 8. *Let Assumptions 1–2 hold, and let Assumption 4 hold. Suppose that at iteration k , Line 14 of Algorithm 1 is reached with $\|g^k\| < \epsilon$ for some $\epsilon > 0$. Let $G^k = \{\nabla M(x^k)\mathbf{pat}^{q_s} : s = 1, \dots, j\}$ be the generator set at iteration k . Additionally, suppose there exists $y \in \mathcal{B}(x^k, \Delta_k)$ so that $\partial f(y) = \mathbf{co} \left\{ \nabla F(y)\mathbf{pat}^{q'_{s'}} : s' = 1, \dots, j' \right\}$ with $\{\mathbf{pat}^{q_1}, \dots, \mathbf{pat}^{q_j}\} \subseteq \{\mathbf{pat}^{q'_1}, \dots, \mathbf{pat}^{q'_{j'}}\}$. Then,*

$$\min_{v \in \partial f(y)} \|v\| < (1 + (\kappa_g + L)\eta_2)\epsilon,$$

where κ_g is as in Assumption 3 and L is as in Assumption 1.

Proof. By Assumption 2, $\eta_2 > 0$. Since Line 14 is reached with $\|g^k\| < \epsilon$, it must be that $\frac{\Delta_k}{\eta_2} \leq \|g^k\| < \epsilon$. By Lemma 1 and the suppositions, there exists $v(g^k) \in \partial f(y)$ so that

$$\|g^k - v(g^k)\| \leq (\kappa_g + L)\Delta_k < (\kappa_g + L)\eta_2\epsilon.$$

Applying the reverse triangle inequality yields

$$\|v(g^k)\| < (1 + (\kappa_g + L)\eta_2)\epsilon,$$

from which the desired result follows. \square

We note that the assumptions that the iteration be acceptable and that y lie within Δ_k of x^k directly tie the result to both the master model gradient $\|g^k\|$ and the trust-region size Δ_k . A termination certificate consisting of these two quantities is analogous to the “optimality certificates” used in [4]. In both cases, the certificate can be interpreted as indicating that two of three necessary conditions are satisfied so that there is a point $y \in \mathcal{B}(x^k, \Delta_k)$ so that an element of $\partial f(y)$ is as small in norm as suggested in Lemma 8. The third necessary condition, which is not as straightforward to check, is that the algorithm’s iterates have become sufficiently clustered around y so that the manifolds active at y are a superset of the manifolds active in $\mathcal{B}(x^k, \Delta_k)$.

We remark here that this dependence on knowing all the manifolds active in a given trust region makes a straightforward analysis of rates based on $\|g_k^f\|$ elusive. Therefore, a comparison of the theoretical worst-case complexity of Algorithm 1 with the rates proven for the nonsmooth methods in [10] is currently elusive.

5. Manifold Sampling as a Stochastic Algorithm. Thus far, no restrictions have been placed on the sample set $Y(x^k, \Delta_k)$, apart from its containment in $\mathcal{B}(x^k, \Delta_k)$. In this section, we consider what happens when $Y(x^k, \Delta_k)$ includes stochastically sampled points, a strategy that results in Section 6 show is fruitful when G^k is built by using Algorithm 3.

If random points are added to $Y(x^k, \Delta_k)$, the sequence of generator sets $\{G^k : k \geq 1\}$ will be a realization of a random variable denoted $\{\tilde{G}^k : k \geq 1\}$. Consequently, the algorithm will be inherently stochastic; the sequence of iterates produced by Algorithm 1 will be random variables $\{\tilde{x}^k : k \geq 1\}$ with realizations $\{x^k : k \geq 1\}$. Similarly, we denote by $\{\tilde{\Delta}_k : k \geq 1\}$, $\{\tilde{s}^k : k \geq 1\}$, and $\{\tilde{g}^k : k \geq 1\}$ sequences of random trust-region radii, trial steps, and master model gradients with respective realizations $\{\Delta_k : k \geq 1\}$, $\{s^k : k \geq 1\}$, and $\{g^k : k \geq 1\}$.

We show in Theorem 9 that the results from Section 4 hold for *any* realization of Algorithm 1. Note that Assumptions 1 and 2 are unaffected by stochasticity in $Y(x^k, \Delta_k)$ (and G^k is similarly unaffected on any iteration). In particular, we note that Assumption 3 ensures that the quality of the component models holds in a deterministic fashion. Assumption 4 has the stochastic analogue of assuming that every iterate in any realization satisfies the deterministic Assumption 4. This is nonrestrictive since Assumption 4 requires that G^k satisfy conditions depending not on $Y(x^k, \Delta_k)$ but only on the sign patterns in $\mathcal{B}(x^k, \Delta_k)$, which is a deterministic set at any iteration.

THEOREM 9. *Let X^* denote the union of all cluster points x^* over all realizations $\{(x^k, \Delta_k, s^k, G^k, g^k) : k \geq 1\}$ in the σ -algebra generated by $\{(\tilde{x}^k, \tilde{\Delta}_k, \tilde{s}^k, \tilde{G}^k, \tilde{g}^k) : k \geq 1\}$. Let Assumptions 1–2 hold, let Assumption 4 hold for each (x^k, Δ_k) in any realization $\{(x^k, \Delta_k) : k \geq 1\}$, and let Algorithm 1 be initialized with $(\tilde{x}^0, \tilde{\Delta}_0) = (x^0, \Delta_0)$. Then, for every $x^* \in X^*$, $0 \in \partial f(x^*)$.*

Proof. The proof follows the same argument as in Section 4.

Let $\{(x^k, \Delta_k, s^k, G^k, g^k) : k \geq 1\}$ be an arbitrary realization of the random sequence $\{(\tilde{x}^k, \tilde{\Delta}_k, \tilde{s}^k, \tilde{G}^k, \tilde{g}^k) : k \geq 1\}$ produced by the stochastic algorithm.

Lemma 1 is independent of the realization, and both the Projection Theorem and Theorem 6 hold independently of Algorithm 1. Lemma 3 holds deterministically for any k where Δ_k and $\|g^k\|$ (produced by G^k) satisfy $\Delta_k < \min \left\{ \frac{1-\eta_1}{\kappa_f}, \eta_2 \right\} \|g^k\|$.

Lemma 4 depends only on f being bounded below (a result of Assumption 1), and thus we get $\Delta_k \rightarrow 0$ for the arbitrary realization.

Lemma 5 holds if $Y(x^k, \Delta_k)$ is stochastic, thereby producing stochastic G^k , because the realization $\{(\Delta_k, g^k) : k \geq 1\}$ having $\liminf_{k \rightarrow \infty} \|g^k\| \neq 0$ would similarly contradict $\Delta_k \rightarrow 0$ (Lemma 4).

We can now prove the theorem. Suppose $\{x^k : k \geq 1\}$ has a cluster point x^* . Then, having proved that all the lemmata hold for the arbitrary realization, a direct application of Theorem 7 to that particular realization gives us that $0 \in \partial f(x^*)$. Since the realization of $\{(\tilde{x}^k, \tilde{\Delta}_k, \tilde{s}^k, \tilde{G}^k, \tilde{g}^k) : k \geq 1\}$ was arbitrary, we have shown the desired result. \square

We note that Theorem 9 is a stronger result than saying that the cluster points of $\{\tilde{x}^k : k \geq 1\}$ almost surely have zero in the subdifferential.

6. Numerical Results. We now examine the performance of variations of the manifold sampling algorithm outlined in Algorithm 1. Throughout this section, we use $x^{(j,p,s)}$ to denote the j th point evaluated on a problem p by an optimization solver s . For derivative-based versions of Algorithm 1, such points correspond solely to the trust-region subproblem solutions (Line 8) and points possibly sampled when constructing the generator set (Line 5); for derivative-free versions, evaluated points may additionally include evaluations performed to ensure that the component models are fully linear in the current trust region (Line 4). We drop the final superscript (s) when the point is the same for all solvers.

6.1. Implementations. In our first tests, we focus on the derivative-free case, when only zero-order information (function values) of F is provided to a solver; we view such solvers as a broader test of the manifold sampling algorithm.

In our implementations of Algorithm 1, the sampling set $Y(x^k, \Delta_k)$ was used for construction of both component models and generator sets (through Algorithm 3). All our implementations employ linear models, m^{F_i} , of each component function F_i . The linear models are constructed so that m^{F_i} interpolates F_i at x^k and is the least-squares regression model for the remainder of the sampling set, $Y(x^k, \Delta_k) \setminus x^k$. At the beginning of iteration k , we set $Y(x^k, \Delta_k)$ to be all points previously evaluated by the algorithm that lie in $\mathcal{B}(x^k, \Delta_k)$. If this results in an underdetermined interpolation (i.e., $\text{rank}(Y(x^k, \Delta_k) - x^k) < n$), then additional points are added to $Y(x^k, \Delta_k)$ as described below.

We tested four variants of Algorithm 1, which differ from one another in how they construct the generator set G^k and how they add points to the sampling set $Y(x^k, \Delta_k)$:

Center Manifold Sampling (CMS): Uses Algorithm 2 to build the generator set G^k ; this generator set does not depend on the sampling set $Y(x^k, \Delta_k)$, which is used solely for constructing the component models. For building these models, the scaled coordinate directions, $\{x^k + \Delta_k e_1, \dots, x^k + \Delta_k e_n\}$ all are added to the sample set $Y(x^k, \Delta_k)$ in cases of underdetermined interpolation.

Greedy Deterministic Manifold Sampling (GDMS): Uses Algorithm 3 to build the generator set G^k . Additional points are not added to the sample set $Y(x^k, \Delta_k)$ unless the linear regression is underdetermined. In the underdetermined case, $n - \text{rank}(Y(x^k, \Delta_k) - x^k)$ directions D in the null space of $Y(x^k, \Delta_k) - x^k$ are generated by means of a (deterministic) QR factorization. After evaluating F along these scaled directions, the associated points $x^k + \Delta_k D$ are added to the sample set $Y(x^k, \Delta_k)$.

Deterministic Manifold Sampling (DMS): Uses Algorithm 3 to build the generator set and adds scaled coordinate directions, $\{x^k + \Delta_k e_1, \dots, x^k + \Delta_k e_n\}$, to the sample set $Y(x^k, \Delta_k)$ every iteration.

Stochastic Manifold Sampling (SMS): Uses Algorithm 3 to build the generator set and adds a set of n points randomly generated from a uniform distribution on $\mathcal{B}(x^k, \Delta_k)$ to the sample set $Y(x^k, \Delta_k)$ every iteration.

The strategy used to add points to the sample set $Y(x^k, \Delta_k)$ in the deterministic variants ensures the full linearity of the models required in Assumption 3. For the stochastic variant SMS, however, the realized sample set $Y(x^k, \Delta_k)$ results in models that do not necessarily satisfy Assumption 3. Consequently, Theorem 9 may not hold for our implementation since such a sample set does not guarantee that the realized models are fully linear (see, e.g., [6, 18]).

In all cases, the quadratic program in Line 7 of Algorithm 1 was solved by the subproblem solver used in [7], which is based on a specialized active set method proposed in [16].

We compared the above variants with a modified version of the minimax method of Grapiglia et al. [11], which we denote GYY. We adjusted the code used in [11] to solve ℓ_1 -problems by changing the nonsmooth subproblem linear program (2.6). We also adjusted stopping tolerances to prevent early termination: we decreased the minimum trust-region radius to 10^{-32} and removed the default criterion of stopping after 10 successive iterations without decrease in the objective.

As a baseline, we also tested two trust-region algorithms for smooth optimization.

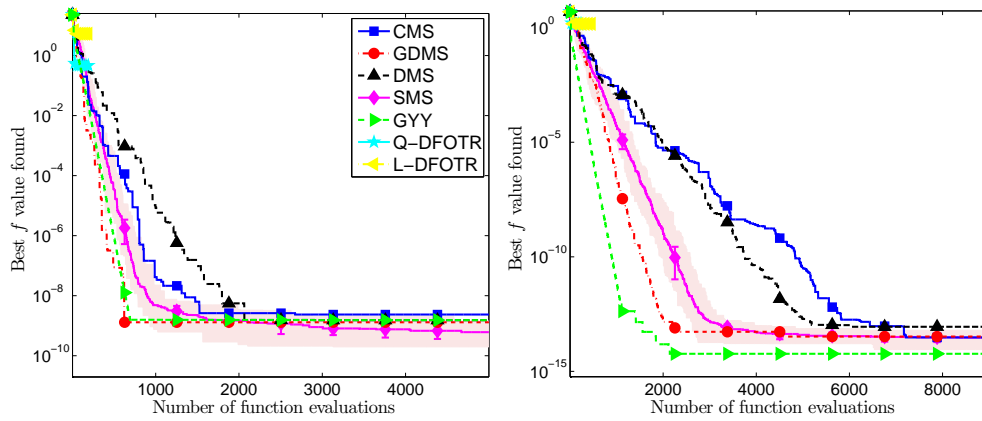


FIGURE 1. Sample trajectories of function values on problem 11 (left), which has $n = r = 4$, and problem 52 (right), which has $n = r = 8$. In SMS, 30 runs were performed: the upper band shows the largest function value obtained at the indicated number of function evaluations and the lower band shows the least function value obtained; the median values are indicated in the line segment connecting the 25th and 75th quantiles of the function values.

The codes L-DFOTR and Q-DFOTR are implementations of the algorithm described in [6] using, respectively, linear and quadratic regression models defined by an appropriately sized, deterministic sample set. These implementations may not converge to a stationary point since they assume a smooth objective function, but they serve as important baselines since they are more efficient at managing their respective sample sets.

By design, all seven of the codes tested employ a trust-region framework, and thus the parameters across the methods can be set equal. The parameter constants were selected to be $\Delta_0 = \max\{1, \|x^{(0,p)}\|_\infty\}$, $\eta_1 = 0.25$, $\eta_2 = 1$, $\gamma_{\text{dec}} = 0.5$, and $\gamma_{\text{inc}} = 2$.

6.2. Test Problems. We consider the ℓ_1 test problems referred to as the “piecewise smooth” test set in [19]. This synthetic test set was selected in part because of the availability of the Jacobian $\nabla F(x)$ for each problem, and thus the subdifferential in (2.4); this is useful for benchmarking purposes. The set is composed of 53 problems of the form (1.1) ranging in dimension from $n = 2$ to $n = 12$, with the number of component functions, r , ranging from n to 65.

A standard starting point, $x^{(0,p)}$, is provided for each problem p in the test set. We note that the objective f is nondifferentiable at $x^{(0,p)}$ for five problems (numbers 9, 10, 29, 30, and 52).

For all test problems, there is neither a guarantee that there is a unique minimizer x^* with $0 \in \partial f(x^*)$ nor a guarantee that $f(x^*) = f(y^*)$ for all x^*, y^* with $0 \in \partial f(x^*) \cap \partial f(y^*)$.

A maximum budget of $1000(n + 1)$ function evaluations was given to each solver for each n -dimensional problem. Solvers were terminated short of this budget only when Δ_k fell below 10^{-32} . We note, however, that for virtually every solver and problem, no successful iterations were found after Δ_k fell below 10^{-18} ; this result is unsurprising given that the experiments were run in double precision.

Figure 1 shows typical trajectories of the best function value found on two problems where $\min f(x) = 0$. The sole stochastic solver (SMS) was run 30 times; Figure 1

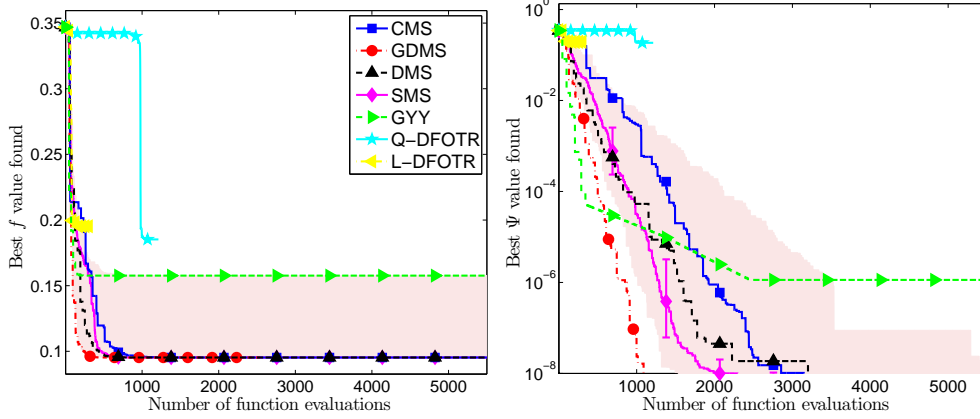


FIGURE 2. *Best function value (left) and stationary measure $\Psi(x)$ (right) found in terms of the number of function evaluations performed for problem 31, a problem with $n = 8$ dimensions and $r = 8$ components. The stationary measure indicates that all instances of SMS find a stationary point, despite the function values associated with these stationary points being different.*

shows that there is little variability across the 30 instances. Figure 2 (left) shows the trajectory on problem 31, where we see that at least one instance of SMS finds a best function value different from that of the majority of SMS instances. In each of these instances, the smooth solvers L-DFOTR and Q-DFOTR struggle to find solutions with function values comparable to those found by the nonsmooth solvers.

6.3. Measuring Stationarity. The behavior seen in Figure 2 (left) suggests that the function values found by a solver may not indicate whether the solver has found a stationary point. We now measure the ability of a solver to identify points close to Clarke stationarity.

Lemma 8 does not guarantee that $(\Delta_k, \|g^k\|)$ provide a measure of stationarity. Instead, we will employ the stationarity measure used for nonsmooth composite optimization in [22] and more recently in [11]. This measure considers the maximum decrease obtained from directional linearizations of f at x ,

$$\Psi(x) = \max_{d: \|d\| \leq 1} f(x) - \|F(x) + \nabla F(x)^T d\|_1. \quad (6.1)$$

From the fact that f is Clarke regular (see, e.g., [5]), a cluster point x^* of $\{x^k\}$ being Clarke stationary is equivalent to the condition $\liminf_{k \rightarrow \infty} \Psi(x^k) = 0$. Such a stationary measure is also readily computed for our benchmark problems since the Jacobian ∇F is known. For example, a variant where the Euclidean norm on d is changed to an ℓ_∞ norm can be obtained by solving the linear optimization problem

$$\min_{d, s} \{e^T s : s \geq F(x^k) + \nabla F(x^k)^T d, s \geq -F(x^k) - \nabla F(x^k)^T d, d \in [-1, 1]^n\}.$$

The importance of using the stationary measure Ψ is highlighted in Figure 2, where we see the performance of seven algorithms on problem 31. Most of the manifold sampling implementations find the same (and largest) amount of function-value decrease, but GYY and an SMS instance converge to a point with a relatively worse function value. Even so, these points all have similar stationary behavior. L-DFOTR and Q-DFOTR fail to find a stationary point.

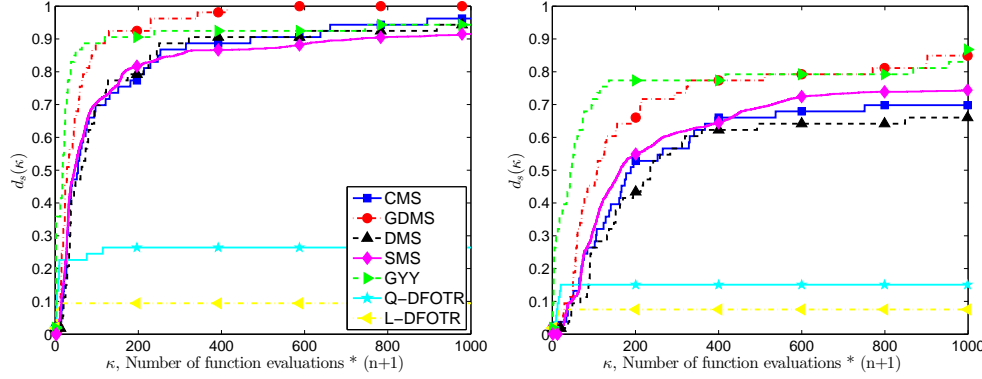


FIGURE 3. Data profiles based on the function value convergence measure (6.2) for $\tau = 10^{-3}$ (left) and $\tau = 10^{-7}$ (right).

6.4. Measuring Performance across the Set. For comparing of performance across the entire test set, we use the data profiles described in [19]. Let S denote the set of solvers we wish to compare, and let P denote the set of test problems. By $t_{p,s}$ we denote the number of function evaluations required for solver $s \in S$ to satisfy a convergence criterion on problem $p \in P$. We use the convention that $t_{p,s} = \infty$ if this convergence criterion is not satisfied within the maximum budget of evaluations.

For $\kappa \geq 0$, the data profile for solver s is then defined by

$$d_s(\kappa) = \frac{1}{|P|} |\{p \in P : t_{p,s} \leq \kappa(n_p + 1)\}|,$$

where n_p is the dimension of problem p .

We first examine the convergence criterion in [19], which is based on the best function value found by an algorithm. In particular, given a tolerance $\tau > 0$, we will say that solver s has converged on problem p when an $x^{(j,p,s)} \in \mathbb{R}^{n_p}$ has been found such that

$$f(x^{(j,p,s)}) \leq f_p + \tau(f(x^{(0,p)}) - f_p), \quad (6.2)$$

where f_p is the least function value obtained across all evaluations of all solvers in S for problem p and where $x^{(0,p)}$ is an initial point common to all solvers. The parameter τ determines how accurate one expects a solution to be in terms of the achievable decrease $f(x^{(0,p)}) - f_p$.

Figure 3 (left) shows that all manifold sampling implementations and GYY successfully find points with function values better than 99.9% of the possible decrease for over 90% of the problems. For the smaller τ , the solvers' performances are more distinguishable. SMS finds decrease at least as good as $(1 - 10^{-7})\%$ of the best-performing method on 75% of the problems, while GDMS and GYY do so for 85% of the problems. The relative success of GDMS over the other deterministic solvers highlights the importance of a judicious use of the function evaluation budget. The quick plateau behavior for the smooth solvers indicate that these solvers are efficient on the problems that they are able to solve.

Our next convergence criterion relates to the stationarity measure (6.1). Given a tolerance $\tau > 0$, we say that convergence has occurred when

$$\Psi(x^{(j,p,s)}) \leq \tau \Psi(x^{(0,p)}). \quad (6.3)$$

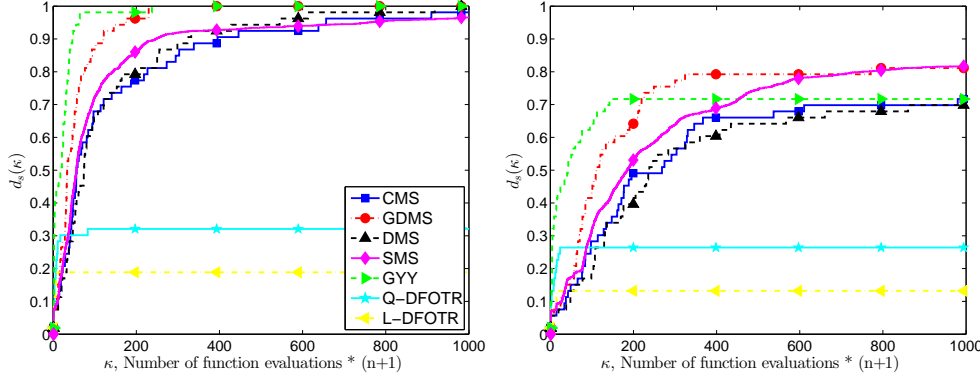


FIGURE 4. Data profiles based on the Ψ convergence measure (6.3) for $\tau = 10^{-3}$ (left) and $\tau = 10^{-7}$ (right).

Using (6.3) to test for convergence in Figure 4, we gain additional insight into the performance of the solvers. Figure 4 (left) shows that Q-DFOTR and L-DFOTR do not find points with Ψ values less than one-thousandth of the stationary measure at $x^{(0,p)}$ on a majority of the benchmark problems, while the other solvers do so for over 95% of the problems. For the more restrictive τ , CMS and DMS perform nearly identically, while SMS and GDMS are shown to be even more robust. GYY is relatively faster at finding small Ψ values in the initial $150(n+1)$ function evaluations. Note that the data profiles for SMS are improved when moving from function value measures (Figure 3) to stationarity measures (Figure 4). We attribute this behavior to the fact that some stochastic instances find stationary points with relatively worse function values (recall Figure 2).

Before proceeding, we note that although these tests show that GDMS is efficient when evaluations are performed sequentially, the other variants have the ability to utilize $n+1$ evaluations concurrently and therefore might prove more useful in a parallel setting.

6.5. Comparison with Gradient Sampling. We also compare the performance between a variant of manifold sampling that uses some gradient information (SMS-G) and GRAD-SAMP, a MATLAB implementation of gradient sampling from [4] that uses gradient information at every evaluated point. This code was run with its default settings and a budget of $1000(n+1)$ Jacobian (and hence gradient) evaluations. Since GRAD-SAMP does not proceed from a nondifferentiable initial point, the five problems with nondifferentiable starting points were perturbed by machine epsilon. We also extended the set of sampling radii in GRAD-SAMP from $\{10^{-4}, 10^{-5}, 10^{-6}\}$ to $\{10^{-4}, 10^{-5}, \dots, 10^{-16}\}$ to avoid early termination.

As suggested by its name, SMS-G is SMS from Section 6.1 with the following modifications. The model building step, Line 4 in Algorithm 1, directly uses the Jacobian $\nabla F(x^k)$ and thus $M(x) = F(x^k) + \nabla F(x^k)(x - x^k)$. Since the default settings in GRAD-SAMP samples $\min(n+10, 2n)$ gradients per iteration, SMS-G has this manifold sampling rate as opposed to the n points sampled each iteration by SMS.

Notice that because SMS-G computes a new Jacobian only immediately following a successful iteration, it incurs at most one Jacobian evaluation per iteration. The manifold sampling step and the evaluation of the trial point in SMS-G require

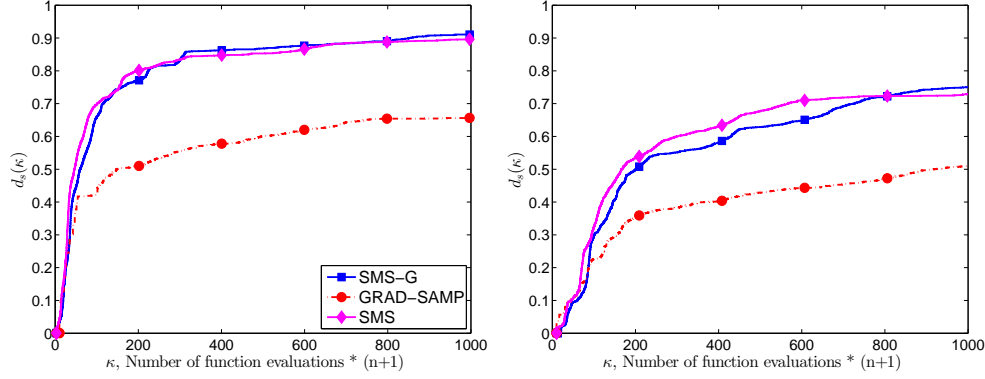


FIGURE 5. Data profiles based on the function value convergence measure (6.2) for $\tau = 10^{-3}$ (left) and $\tau = 10^{-7}$ (right). Note that SMS samples n points per iteration while SMS-G and GRAD-SAMP sample $\min\{n + 10, 2n\}$ points per iteration.

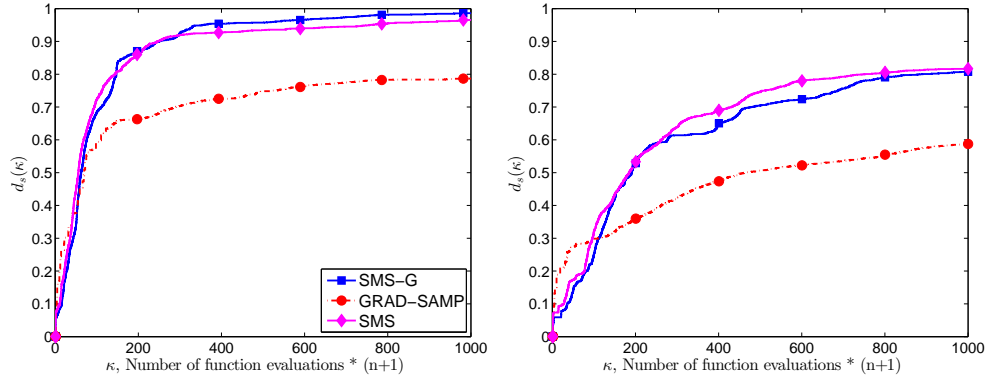


FIGURE 6. Data profiles based on the Ψ convergence measure (6.3) for $\tau = 10^{-3}$ (left) and $\tau = 10^{-7}$ (right).

only function evaluations (i.e., not Jacobian evaluations), and so SMS-G incurs at most $\min(n + 11, 2n + 1)$ function evaluations per iteration. On the other hand, GRAD-SAMP can require a bundle of $\min(n + 11, 2n + 1)$ gradient (and hence Jacobian) and corresponding function evaluations per iteration. Thus, in our data profiles measured in terms of function evaluations, every function evaluation used by GRAD-SAMP entails a Jacobian evaluation; SMS-G function evaluations include a Jacobian only evaluation for a fairly small (always less than 20%) proportion of the function evaluations.

Data profiles are shown in Figure 5 and Figure 6 for an experiment where 30 stochastic runs were performed for both solvers. We also compare results with 30 instances of the Jacobian-free SMS described in Section 6.1. The gradient sampling method performs significantly worse than either manifold sampling method. Furthermore, the similar performance exhibited by the Jacobian-based SMS-G and the Jacobian-free SMS indicates that the performance of SMS-G would likely further improve if the sampling rate were reduced.

7. Discussion. The driving force behind the proposed manifold sampling algorithm is that search directions are computed by using a finitely generated set,

$$\text{co} \left\{ \nabla M(x^k) \text{sign}(y^j) + \sum_{i: [\text{sign}(y^j)]_i = 0} t_i \nabla m^{F_i}(x^k) : t \in \{-1, 0, 1\}^r, y^j \in Y(x^k; \Delta_k) \right\},$$

which differs from the finitely generated set used by gradient sampling,

$$\text{co} \{ \nabla M(y^j) \text{sign}(y^j) : y^j \in \mathcal{D} \cap Y(x^k; \Delta_k) \}.$$

Our tests on ℓ_1 functions show that the manifold sampling strategies compare favorably with a gradient sampling approach.

Although the present work targeted composite problems (1.2) for the particular nonsmooth function $h(u) = \|u\|_1$, the approach can be extended to other functions h , provided one can classify points into smooth manifolds (either with zero-order information or inexact first-order information). In the case of ℓ_1 functions, this classification was determined by the trivial evaluation of the sign pattern of $F(y)$ for a sample $y \in Y(x, \Delta)$. In the case of minimax objective functions of the form

$$h(u) = \max_{i=1, \dots, r} u_i \quad \text{or} \quad \max_{i=1, \dots, r} |u_i|,$$

this classification is determined by what Hare and Nutini refer to as the “active set” at a point [13]. In general, in any setting where the form of the subdifferential ∂h at any point is known, a setting that subsumes much of the work in the nonsmooth composite optimization literature, an analogous version of Algorithm 1 can be proposed.

In particular, the manifold sampling approach does not rely on convexity of the function h . This is in contrast to methods that solve the nonsmooth subproblem (2.6). An example of such a method is the GYY code modified from [11], which we showed can slightly outperform manifold sampling on ℓ_1 problems.

In addition to the extensions above, we see opportunities for employing approximate second-order information to accelerate convergence. In this paper, we proved results only for a certain linear master model, which may be used as a baseline for proving results about search directions that use curvature information about the component functions F_i . Using second-order information in the derivative form of a gradient sampling algorithm through L-BFGS steps was explored in [7]. Curvature information in derivative-free nonsmooth optimization was also explored in [14].

We are also interested in efficient and greedy updates of sample sets for manifolds and/or models in both settings where function (and Jacobian) evaluations are performed sequentially and concurrently. Our implementation of GDMS is a first step in this direction. Furthermore, natural questions arise about the tradeoff between richness of manifold information to reach early termination of the inner while loop in Algorithm 1 and efficiency to guarantee that the sample set is, for example, well poised for model building.

Acknowledgments. This material is based upon work supported by the U.S. Department of Energy, Office of Science, under Contract DE-AC02-06CH11357. We are grateful to Katya Scheinberg for the code we refer to as DFOTR, Geovani Grapiglia for the GYY implementation, and to Frank Curtis and Xiaocun Que for the use of their specialized active set method in solving our minimum norm element QP subproblems. Wild is grateful to Aswin Kannan for early numerical experiments testing methods for composite black-box optimization.

REFERENCES

- [1] CHARLES AUDET AND JOHN E. DENNIS, JR., *Mesh adaptive direct search algorithms for constrained optimization*, SIAM Journal on Optimization, 17 (2006), pp. 188–217. (Cited on page 4.)
- [2] DIMITRI P. BERTSEKAS, *Convex Optimization Theory*, Athena Scientific, 2009. (Cited on page 9.)
- [3] JAMES V. BURKE, *Descent methods for composite nondifferentiable optimization problems*, Mathematical Programming, 33 (1985), pp. 260–279. (Cited on page 4.)
- [4] J. V. BURKE, A. S. LEWIS, AND M. L. OVERTON, *A robust gradient sampling algorithm for nonsmooth, nonconvex optimization*, SIAM Journal on Optimization, 15 (2005), pp. 751–779. (Cited on pages 4, 13, and 19.)
- [5] FRANK H. CLARKE, *Optimization and Nonsmooth Analysis*, John Wiley & Sons, 1983. (Cited on pages 2, 3, and 17.)
- [6] ANDREW R. CONN, KATYA SCHEINBERG, AND LUÍS N. VICENTE, *Introduction to Derivative-Free Optimization*, MPS/SIAM Series on Optimization, SIAM, Philadelphia, PA, 2009. (Cited on pages 6, 15, and 16.)
- [7] FRANK E. CURTIS AND XIAOCUN QUE, *An adaptive gradient sampling algorithm for non-smooth optimization*, Optimization Methods and Software, 28 (2013), pp. 1302–1324. (Cited on pages 15 and 21.)
- [8] R. FLETCHER, *A model algorithm for composite nondifferentiable optimization problems*, in Nondifferential and Variational Techniques in Optimization, D. C. Sorensen and R. J.-B. Wets, eds., vol. 17 of Mathematical Programming Studies, Springer Berlin Heidelberg, 1982, pp. 67–76. (Cited on page 4.)
- [9] ROGER FLETCHER, *Practical Methods of Optimization*, John Wiley & Sons, New York, second ed., 1987. (Cited on page 4.)
- [10] R. GARMANJANI, D. JÚDICE, AND L. N. VICENTE, *Trust-region methods without using derivatives: Worst case complexity and the non-smooth case*, Preprint 15-03, Dept. Mathematics, University of Coimbra, March 2015. (Cited on pages 4 and 13.)
- [11] GEOVANI NUNES GRAPIGLIA, JINYUN YUAN, AND YA-XIANG YUAN, *A derivative-free trust-region algorithm for composite nonsmooth optimization*, Computational and Applied Mathematics, (2014), pp. 1–25. (Cited on pages 4, 15, 17, and 21.)
- [12] ANDREAS GRIEWANK, ANDREA WALTHER, SABRINA FIEGE, AND TORSTEN BOSSE, *On Lipschitz optimization based on gray-box piecewise linearization*, Mathematical Programming, (2015), pp. 1–33. (Cited on page 4.)
- [13] W. HARE AND J. NUTINI, *A derivative-free approximate gradient sampling algorithm for finite minimax problems*, Computational Optimization and Applications, 56 (2013), pp. 1–38. (Cited on pages 4, 13, and 21.)
- [14] NAPSU KARMTSA AND A. M. BAGIROV, *Limited memory discrete gradient bundle method for nonsmooth derivative-free optimization*, Optimization, 61 (2012), pp. 1491–1509. (Cited on page 21.)
- [15] KRZYSZTOF C. KIWIŁ, *Methods of Descent for Nondifferentiable Optimization*, vol. 1133 of Lecture Notes in Mathematics, Springer Berlin Heidelberg, 1985. (Cited on page 3.)
- [16] ———, *A method for solving certian quadratic programming problems arising in nonsmooth optimization*, IMA Journal of Numerical Analysis, (1986), pp. 137–152. (Cited on page 15.)
- [17] ———, *A nonderivative version of the gradient sampling algorithm for nonsmooth nonconvex optimization*, SIAM Journal on Optimization, 20 (2010), pp. 1983–1994. (Cited on pages 4 and 5.)
- [18] JEFFREY LARSON AND STEPHEN BILLUPS, *Stochastic derivative-free optimization using a trust region framework*, preprint, 2015. http://www.optimization-online.org/DB_FILE/2013/07/3978. (Cited on page 15.)
- [19] JORGE J. MORÉ AND STEFAN M. WILD, *Benchmarking derivative-free optimization algorithms*, SIAM Journal on Optimization, 20 (2009), pp. 172–191. (Cited on pages 16 and 18.)
- [20] DOBRIVOJE POPOVIC AND ANDREW R. TEEL, *Direct search methods for nonsmooth optimization*, in Proceedings of the 43rd IEEE Conference on Decision and Control, 2004. (Cited on page 4.)
- [21] R. TYRRELL ROCKAFELLAR AND ROGER J.-B. WETS, *Variational Analysis*, vol. 317, Springer, 2009. (Cited on page 11.)
- [22] YA-XIANG YUAN, *Conditions for convergence of trust region algorithms for nonsmooth optimization*, Mathematical Programming, (1985), pp. 220–228. (Cited on pages 4 and 17.)

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory (“Argonne”). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.